

Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) EP 0 977 402 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:  
02.02.2000 Bulletin 2000/05

(51) Int. Cl.<sup>7</sup>: H04L 12/56, H04Q 11/04

(21) Application number: 99303877.7

(22) Date of filing: 23.04.1999

(84) Designated Contracting States:  
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE  
Designated Extension States:  
AL LT LV MK RO SI

(72) Inventors:  
• Duckering, Brent G.  
Vancouver, British Columbia V5V 4H4 (CA)  
• Robotham, Rob E.  
Ottawa, Ontario K1H 5L8 (CA)

(30) Priority: 23.04.1998 US 64623

(74) Representative:  
Hutchins, Michael Richard et al  
FRY HEATH & SPENCE  
The Old College  
53 High Street  
Horley Surrey RH6 7BN (GB)

(71) Applicant:  
Newbridge Networks Corporation  
Kanata, Ontario K2K 2E6 (CA)

(54) Fair share scheduling of multiple service classes with prioritized shaping

(57) An apparatus for scheduling multi-service category ATM cell traffic through contention points in the ATM network is provided. The service categories have predefined delivery priorities according to quality of service guarantees. To satisfy these priorities while maintaining fair treatment to low priority connections, aging markers are incrementally assigned to queued cells and these markers in combination with priority data are used to determine which connection is serviced next. A method of scheduling traffic through the network is also disclosed.




FIGURE 2

EP 0 977 402 A2

AL

## Description

### Field of the Invention

[0001] This invention relates to systems and methods for fairly arbitrating between multiple scheduled events and more particularly to systems and methods for the fair-share scheduling of multiple service categories at queuing or congestion points in an ATM network.

### Background of the Invention

[0002] Asynchronous transfer mode (ATM) is gaining rapid recognition as the technology of choice for the transmission of broadband information. The fixed length cell size employed in ATM technology supports the delivery of a wide range of multimedia information including speech, video and data. In order to service this wide variety of information the ATM Traffic Management Forum has specified that ATM traffic be divided into five service categories. Two of these service categories are for real-time transmission, namely constant bit rate (CBR) and real-time variable bit rate (rtVBR). There are, as well, three non-real-time categories, namely non-real-time variable bit rate (nrt-VBR); unspecified bit rate (UBR) and available bit rate (ABR).

[0003] As the various service categories are designed to carry different types of information, different Quality of Service (QoS) parameters apply. The ATM Traffic Management Forum has listed six Quality of Service parameters namely: peak-to-peak Cell Delay Variation (peak-to-peak CDV); maximum Cell Transfer Delay (Max CTD); Cell Loss Ratio (CLR); Cell Error Ratio (CER); Severely Errored Cell Block Ratio (SECBR) and Cell Misinsertion Ratio (CMR). Of these six, two are of greatest concern when scheduling. These are the maximum Cell Transfer Delay (Max CTD) and Cell Delay Variation (CDV). Traffic of different service classes, and even traffic of the same service class may have different delay requirements. Fair scheduling between connections of the various ATM service categories onto a single resource while meeting the multiple Quality of Service delay guarantees is difficult to achieve.

[0004] One prior art method for scheduling such traffic, known as shaping, involves calculating a theoretical emission time (TET) for each active connection and attempting to service each connection as close to that theoretical emission time as possible. Unfortunately, most connections cannot be serviced exactly at their TET due to finite output time requirements for the shared resource, as well as scheduling collisions where multiple connections are scheduled to output at the same time. Though the finite output times cannot be avoided, the effects of scheduling collisions can and should be reduced through prioritization. Without prioritization, a connection with loose delay requirements may actually get serviced prior to a connection with tight delay requirements scheduled at the same TET.

[0005] One proposed solution to this problem is to use two or more shaping devices (or shapers) to schedule the connections where connections with similar Quality of Service delay requirements are scheduled in the same shaper. Exhaustive servicing is commonly used to determine which shaper gets to send a cell at any given time. Prioritized shaping is discussed in US application S.N. which is based on provisional application 60/020,642 filed June 27, 1996. The contents of the aforementioned application are incorporated herein by reference. When more than one calendar (as defined in the aforementioned application) contains a connection ready to send a cell, only the highest priority calendar gets to transmit. All lower priority calendars have to wait until all calendars with higher priority have nothing ready to send. This calendar prioritization can be based upon the CDV and possibly max CTD requirements of the traffic scheduled in each calendar. There are, however, certain limitations which can be identified in connection with the aforementioned multiple shaper or calendar solution. For example, there is no quantitative link between the CDV and max CTD requirements between each shaper's traffic and the prioritization of that shaper respective of all other shapers. When a high priority shaper schedules several connections over a relatively short transient period, lower priority shapers have to wait until this transient finishes before they can resume sending cells from their connections. During these transients, low priority connections may violate their delay requirements while those scheduled on high priority shapers may actually beat their delay requirements. Thus, high priority connections may actually better their QoS at the expense of lower priority connections. To properly support QoS delay requirements, shaper priorities must be dynamically assigned based upon the amount of time that their connections have been delayed from ideal theoretical emission times.

[0006] Further, heretofore traffic is either shaped or placed in work-conserving, weighted-fair-queues. Shaping works well for real-time traffic where rates are explicitly guaranteed and weighted-fair-queueing works well for non-real-time traffic that does not need to be prioritized or rate limited. However, most non-real-time traffic has an associated minimum rate that must be guaranteed and a maximum rate that cannot be exceeded. It is also often desirable to provide work conservation between these two limits to ensure full resource usage as precise instantaneous bandwidth availability is rarely known for non-real-time traffic. An example would be non-real-time variable bit rate where the minimum guaranteed scheduled rate is sustainable cell rate (SCR) though output rate will be lower when an upstream source sends below SCR. The maximum rate is peak cell rate (PCR) with an associated burst tolerance (BT) that must be respected. In this case it is not desirable to simply shape up to PCR as this may infringe upon rate guarantees of the available bit rate

(ABR) and possibly unspecified bit rate (UBR) connections. On the other hand, it is also not desirable to simply shape to SCR as this does not allow the non-real-time variable bit rate connection to make use of any additional unused bandwidth. This minimum rate guarantee with work conservation up to a maximum rate cannot be provided for by shaping or weighted-fair-queuing alone.

[0007] Furthermore, scheduling using a single device, namely a shaper or a weighted-fair-queue, does not provide proper support for rate-based backpressure where all non-real-time connections are throttled back to minimum required rates. Weighted-fair-queuing by itself is not able to throttle back all connections due to its work-conserving nature, while shaping is far too slow to respond. For a shaper to throttle back the rates of all connections scheduled, these connections would either have to be descheduled and rescheduled with a longer period or the shaper would have to wait for each connection to be serviced before rescheduling it with a longer period. At the conclusion of the backpressure, the same procedure would have to be followed for increasing the rates of the connections. As thousands of connections may be scheduled these procedures would take a very long time to complete. The net effect is an increased delay in the control loop which is therefore harder to stabilize.

#### Summary of the Invention

[0008] It is therefore an object of the present invention to provide fair share scheduling of multiple scheduled events of different service classes with exhaustive by age priority servicing.

[0009] It is a further object of the present invention to provide fair share scheduling with work conservation and prioritized minimum rate.

[0010] It is also an object of the present invention to provide fair share scheduling and peak rate limiting.

[0011] Therefore in accordance with a first aspect of the present invention there is provided an apparatus for arbitrating service between contending multiple scheduled events each having a predefined service delivery priority and a preset delay factor, where the delay factor corresponds directly to the CDV, and also relates to max CTD. Note that max CTD is also a function of the variance in the input stream since shapers use delay to smooth out this variance. The apparatus comprises queuing means to place sequential events in per sequence queues; scheduling means to determine sequence service times and to group sequences of similar service quality; aging means to monitor sequence groups and to record aging components for each group, each aging component representing the largest delay suffered by the sequences in the group; and servicing means to selectively service a sequence having the highest aggregate of service priority and aging component.

[0012] In accordance with a second aspect of the present invention there is provided an apparatus for arbitrating between service connections at a contention point in an ATM network, the service connections including multiple service categories each with predefined Quality of Service guarantees. The apparatus comprises queuing means for arranging sequential items in per sequence queues; scheduling means to determine sequence service times and to group sequences of similar service quality; aging means to monitor sequence groups and to introduce aging markers to sequence groups that have been waiting the longest to send; and servicing means to selectively service a sequence having the highest aggregate of service priority and aging markers.

[0013] In a preferred embodiment of this aspect of the invention the scheduling means includes parallel shaper and weighted-fair-queuing means for specific service categories whereby the shaper provides a minimum rate guarantee and the weighted-fair-queuing means provides optimum bandwidth usage.

[0014] In accordance with a further embodiment of this aspect of the invention shaping means may also be provided to limit peak rate offered to sequences by the weighted-fair-queuing means.

[0015] In accordance with a third aspect of the present invention there is provided a method of fairly scheduling sequential events of various service categories at a queuing point in an ATM network in order to satisfy Quality of Service guarantees. The method comprises the steps of placing sequential events into per-sequence queues; prioritizing the per-sequence queues in sequence groups of similar service quality; recording a timing marker for each sequence group, the timing marker representing the largest delay suffered by the sequences in the group; accumulating priority and timer marker data respecting each group; and selectively servicing the sequence having the highest aggregate of service priority and timing markers.

#### Brief Description of the Drawings

[0016] The invention will now be described in greater detail with reference to the attached drawings wherein:

Figure 1 is a sample multiple calendar scheduler according to the prior art;

Figure 2 illustrates exhaustive by age priority servicing according to the present invention; and

Figure 3 shows work conservation with prioritized minimum rate and rate limiting.

#### Detailed Description of the Invention

[0017] Figure 1 demonstrates an implementation of a prior art queue servicing solution using multiple shaping calendars 12 which constitutes a particular form of a traffic shaper. The calendars 12 are prioritized. When

more than one calendar 12 contains a connection ready to send a cell only the highest priority calendar gets to transmit. All lower priority calendars have to wait until all calendars with higher priority have nothing ready to send. In Figure 1 a low priority weighted-fair-queuing (WFQ) scheduler 14 which could be in the form of a calendar, a tag comparator or other WFQ method has been added to provide work-conserving support for low priority connections such as UBR. Calendars could be divided based upon the traffic type (real-time or non-real-time) service category (CBR, rtVBR, nrt-VBR, ABR, UBR) delay requirements (such as CDV and max CTD) requirements or a combination of these such as delay requirements for real-time traffic and service category for non-real-time traffic. Since Quality of Service (QoS) can be related to network node CDV and max CTD requirements, dividing calendars within the node based upon these requirements prioritize traffic based upon QoS.

**[0018]** Generally speaking, traffic with tight CDV requirements should be transmitted as close to the theoretical emission time as possible. However, if several connections with tight CDV are scheduled in proximity to each other and receive exhaustive priority, it is possible that a connection with a larger CDV scheduled near the beginning of the high priority group may get delayed beyond its CDV requirement.

**[0019]** To avoid this, in accordance with the present invention, an aging factor is added which allows low priority large CDV connections to boost their priority as they get delayed close to the CDV limit for the node. An exhaustive by age priority servicing method is used to service between the multiple shapers. This will also limit max CTD for connections on lower priority shapers, though quantizing max CTD for shaped traffic requires knowledge of the traffic pattern at the input of the shaper.

**[0020]** In the following description reference is made to calendars. As used herein a calendar is a traffic shaping entity which, in effect, defines a succession of time intervals, in each of which are identified queues to be serviced during that interval.

**[0021]** Figure 2 demonstrates one implementation of exhaustive by age priority servicing. Connections are shaped using multiple calendars 20, where each calendar has a programmable 1-point CDV associated with it. The 1-point CDV is described in ATM Forum's Traffic Management Specification Version 4.0. In the multiple calendar architecture illustrated in Figure 2 the position of each queue within the calendar is determined by the Generic Cell Rate Algorithm (GCRA) also defined in the aforementioned ATM Forum's Traffic Management Specification. In this example a WFQ calendar 22 is also available to provide work conservation for low priority connections, though no aging is associated with it. One or more age thresholds are provided for each shaping calendar to support aging as the calendar gets behind in service relative to TET. One threshold could

be set at or near the desired CDV limit for the calendar, while any others could be set above and/or below the CDV to provide additional resolution. An aging bitmap 24 is associated with the group of calendars, with one bit 26 per threshold per calendar, plus one additional bit per calendar to indicate that the calendar has a connection ready to be serviced.

**[0022]** When a calendar reaches the time where a connection is scheduled to transmit, the first bit in the age bitmap is set to indicate that a connection is ready. If the connection cannot be serviced due to other connections waiting in higher priority calendars, additional bits will be set in the bitmap as each aging threshold is crossed. When determining which connection should be serviced next, the aging bitmap is read and the earliest scheduled connection on the highest priority calendar with the highest count of age bits set is serviced. This assumes age takes precedence to priority. At this point the age bitmap is updated if necessary to reflect the servicing.

**[0023]** This example demonstrates the use of a two pointer shaper for determining how far behind each calendar is. The Real Time Pointer (RTP) points to the ideal service point, while the Virtual Time Pointer (VTP) points to the actual service point. The RTP constantly increments at fixed intervals, whereas the VTP points to a particular time slot in the calendar which defines the next connection, i.e. cell to be serviced. The VTP can only be incremented if a connection is serviced from the calendar or if the RTP increments and there are no connections ready to be serviced. The distance between the two pointers determines currently how delayed the oldest connection is from its TET, and can be calculated every time the RTP increments.

**[0024]** It would be possible to use this aging mechanism for other implementations of shapers as well. For example, if connections are removed from the calendar as soon as their TET is reached and placed in a queue associated with the calendar until they can be serviced, their TET values can be used to set the aging bits. In this case, the connection at the head of the queue is the oldest, and its TET value can be compared to the current time and aging thresholds in order to set the aging bits. Furthermore, the shaper does not have to take the form of a calendar.

**[0025]** Although this example demonstrates the use of exhaustive by age priority servicing as it relates to shaping ATM connections, the same principles can be used for any application where multiple scheduled events have to be arbitrated between based upon priority and delay variances.

**[0026]** This method provides a much better means of ensuring that CDV requirements are not violated during scheduling collisions, as compared to simple exhaustive servicing. Large scheduling collisions in high priority calendars will not be able to starve lower priority connections. CDV limits for all priorities can be guaranteed assuming the scheduler is not overbooked, unlike

exhaustive servicing that can only guarantee CDV limits to higher priority connections when the scheduler is more than minimally booked.

[0027] Implementing additional threshold levels below network node CDV requirements provides a better means to distribute the effects of a minor scheduling collision among all priorities, as opposed to only involving low priority traffic until CDV limits have been reached. Implementing additional threshold levels above CDV can minimize the relative size of CDV violations by distributing violations across all priorities during excessive scheduling collisions, instead of violating only low priority connections. Of course for some traffic classes it may be more desirable to avoid CDV violations for higher priority traffic at the expense of large CDV violations for lower priority traffic during large collisions, instead of spreading the violations equally. In this case aging threshold levels above CDV may not be acceptable. Thresholds can also be set to differentiate between calendars shaping non-real-time traffic, although these thresholds should be large enough to prevent non real-time traffic from ever obtaining higher priority than real-time traffic.

[0028] By scheduling a connection with both a shaper and a WFQ at the same time, it is possible to provide a minimum rate guarantee with work conservation providing any additionally available bandwidth. The shaper will shape to the minimum rate, while the WFQ provides everything above. This supports rate-based backpressure, where the non-real-time traffic can be throttled back to minimum rates simply by halting service of the WFQ.

[0029] Figure 3 demonstrates a sample configuration for work conservation with prioritized minimum rate and peak rate limiting. In this example the nrt-VBR is shaped at SCR in a separate calendar 30 from the calendar 32 for ABR and UBR. This gives the nrt-VBR priority during transients where insufficient bandwidth is available to provide minimum rates to all three non-real time classes. Aging can be utilized to prevent the ABR and UBR from being starved completely during the transients. ABR and UBR are shaped to their MCR values. In Figure 3 a separate shaper is provided for peak rate limiting as will be discussed later. WFQ weights are set to provide a fair-share of the remaining bandwidth to all of the non-real-time connections. This example assumes that fair-share is relative to the difference between the peak rate and the minimum rate. If the fair-share calculations for ABR explicit rate (ER) calculations differs from this, the WFQ weights might have to be modified accordingly.

[0030] When scheduling on two calendars 30, 36 simultaneously, care must be taken to ensure that cell sequence integrity is maintained, and that empty cells are not inadvertently sent out due to connections with only one cell queued being scheduled on two calendars. Cell sequence integrity can easily be preserved by scheduling connections rather than cells. This ensures

that the cell at the head of the queue is first to go out regardless as to which of the two calendars is first to service the connection. There is always a possibility that a connection with only one cell queued can be scheduled in two calendars, though the probability of this can be substantially reduced. When a cell arrives to an empty queue, the connection need only be scheduled in the shaper. If the shaper sends this cell out before the next cell arrives, then the connection again only needs to be scheduled in the shaper since the cells are arriving at or below the minimum rate. However, if another cell arrives prior to the first being sent, the connection should be scheduled into the WFQ as well. After this, the connection will be rescheduled every time it is serviced by either calendar until the connection's queue empties. When a queue that is scheduled in two calendars empties, the connection will have to be removed from the calendar in which it is still scheduled. This descheduling will have to occur prior to the next output cell time to ensure that the empty queue is not accidentally serviced.

[0031] As WFQ is work conserving, there may be a need to limit peak rates. Peak rate limiting can be achieved by rescheduling, with a shaper 34, connections that have been serviced too early by the WFQ 36. To determine that a connection has been serviced too early by the WFQ 36, a theoretical emission time can be calculated based upon the peak rate at each emission time, and compared to the actual emission time (AET). If the AET is less than the TET, the connection has been serviced too early and will have to be scheduled in the shaper 34 at an emission time of TET offset by the period corresponding to the peak rate. After the connection has been serviced by the shaper, the average rate over two cell times is no longer in violation of the peak rate, so the connection can again be scheduled in the WFQ 36, assuming more cells are available. This provides for a maximum CDV of one peak rate cell time in addition to the maximum CDV provided by the shaper, and ensures that the peak rate is never violated over a two cell emission interval.

[0032] For the sample configuration of Figure 3, shaper n (34) provides for peak rate limiting the output of the WFQ 36. Depending on design limitations, it may be possible to use one of the minimum rate shapers to also provide for peak rate limiting if it is possible to schedule the same connection twice within the same calendar. If not, the separate peak rate shaper will only get utilized when the WFQ is sufficiently underbooked.

[0033] In this sample configuration, the peak rate for nrt-VBR is PCR-SCR, but this must also respect the burst tolerance. To ensure nrt-VBR burst tolerance (BT) is not exceeded, credits ( $\tau$ ) can be accumulated whenever the connection sends at less than SCR, up to a limit of BT. Only if there are enough credits accumulated for the transmission of at least one cell at PCR will the VBR connection be eligible for scheduling in the WFQ calendar. Credits are subtracted every time the connec-

tion is serviced by the WFQ calendar. As soon as there are insufficient credits, the connection will no longer be eligible for WFQ scheduling until its queue runs empty, as a non-empty queue indicates that the connection is still being shaped to SCR.

[0034] By rescheduling connections that have reached their peak rates in a shaper, connections can be capped at their peak rates independently of their weights in the WFQ. This independence from WFQ weights allows connections that have small peak rates relative to their WFQ weights to be capped without interfering with the work-conserving service of those connections with relatively high peak rates.

[0035] Although specific embodiments of the invention have been described and illustrated it will be apparent to one skilled in the art that various modifications can be made to these embodiments without departing from the basic inventive concepts. It is to be understood that such modifications, to the extent possible, fall within the scope of the invention as defined by the appended claims.

#### Claims

1. An apparatus for arbitrating service between contending multiple scheduled events having a predefined service delivery priority and a preset delay factor, said apparatus comprising:

queuing means to place sequential items in per sequence queues;

scheduling means to determine sequence service times and to group sequences of similar service quality;

aging means to monitor sequence groups and to record aging components for each group, each of said aging components representing the largest delay suffered by sequences in said group; and

servicing means to selectively service a sequence having the highest aggregate of service priority and aging component.

2. An apparatus as defined in claim 1 wherein said delay factor corresponds directly to Cell Delay Variation (CDV).
3. An apparatus as defined in claim 1 wherein said delay factor relates to maximum Cell Transfer Delay (maxCTD).
4. An apparatus for arbitrating between service connections at a contention point in an ATM network, said service connections including multiple service categories each with predefined quality of service guarantees, said apparatus comprising:

queuing means for arranging sequential items

in per sequence queues;

scheduling means to determine sequence service times and to group sequences of similar service quality;

aging means to monitor sequence groups and to introduce aging markers to sequence groups that have been waiting beyond set thresholds to send; and

servicing means to selectively service a sequence having the highest aggregate of service priorities and aging markers.

5. An apparatus as defined in claim 1 wherein said scheduling means includes multiple shaping means.
6. An apparatus as defined in claim 5 further including one or more work-conserving means.
7. An apparatus as defined in claim 5 wherein said shaping means are shaping calendars.
8. An apparatus as defined in claim 6 wherein said work-conserving means is a weighted-fair-queuing scheduler.
9. An apparatus as defined in claim 4 wherein said aging means introduces multiple aging markers to each sequence group, an aging marker being added each time a time interval is passed without said connection being serviced.
10. An apparatus as defined in claim 4 wherein said scheduling means includes parallel shaper and weighted fair queuing means for specific service categories whereby said shaper provides a minimum rate guarantee and said weighted fair queuing means provides optimum bandwidth usage.
11. An apparatus as defined in claim 10 wherein said parallel shaping means and said weighted fair queuing means supports rate controlled backpressure for specific service categories.
12. An apparatus as defined in claim 10 wherein said specific service categories may be scheduled in either a shaper or a weighted fair queuing means in order to limit the peak cell delivery rate.
13. A method of fairly scheduling sequential events of various service categories at a queuing point in an ATM network in order to satisfy Quality of Service guarantees, the method comprising:

placing sequential events into per sequence queues;

prioritizing said per sequence queues in sequence groups of similar service quality;

recording a timing marker for each sequence group, said timing marker representing the largest delay suffered by said sequence queue in said sequence group;

accumulating priority and timing marker data 5  
respecting each group; and

selectively servicing the sequence queue having the highest aggregate of priority and timing markers.

10

15

20

25

30

35

40

45

50

55




FIGURE 1

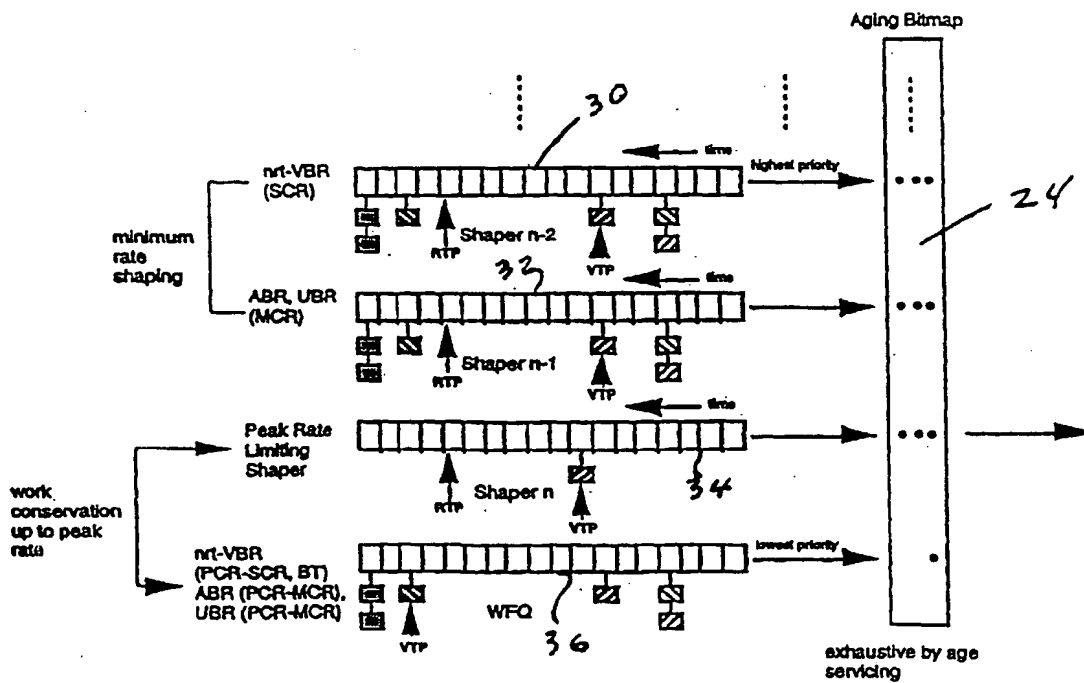


FIGURE 3



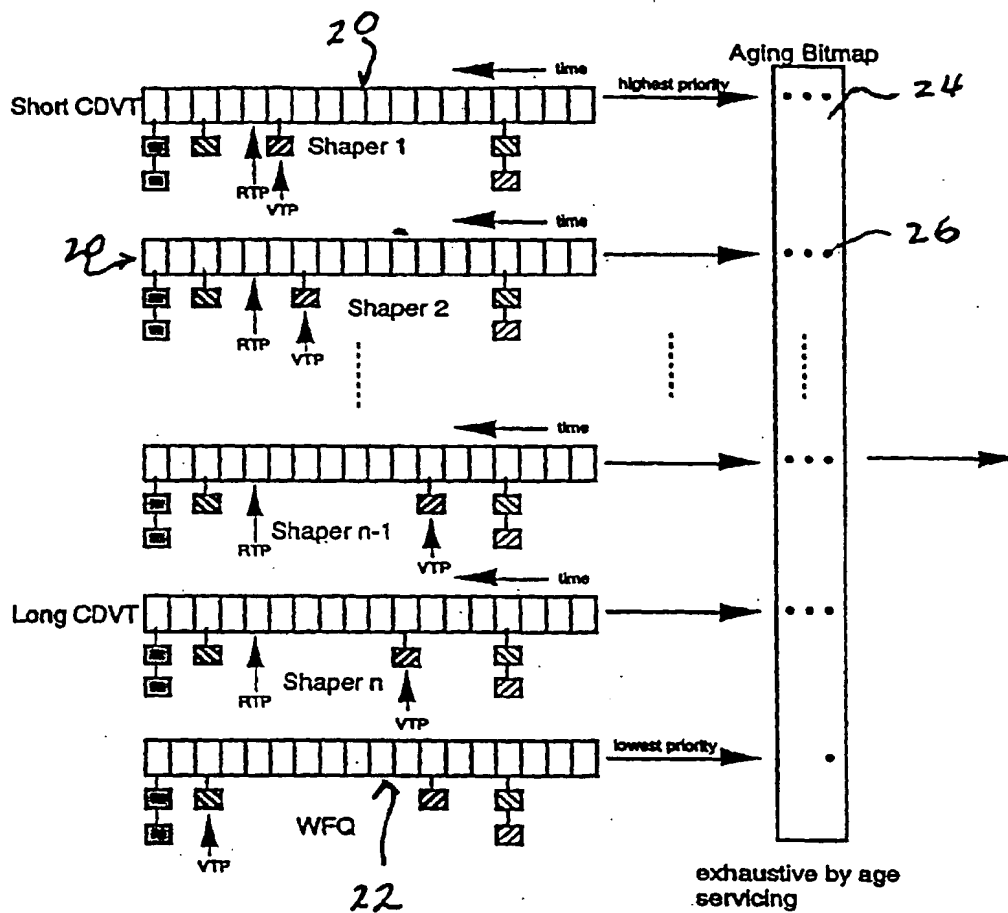


FIGURE 2

**THIS PAGE BLANK (USPTO)**




FIGURE 1




FIGURE 3





FIGURE 2



Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) **EP 0 977 402 A3**

(12) **EUROPEAN PATENT APPLICATION**

(88) Date of publication A3:  
10.01.2001 Bulletin 2001/02

(51) Int. Cl.<sup>7</sup>: **H04L 12/56**, H04Q 11/04

(43) Date of publication A2:  
02.02.2000 Bulletin 2000/05

(21) Application number: **99303877.7**

(22) Date of filing: **23.04.1999**

(84) Designated Contracting States:  
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE**  
Designated Extension States:  
**AL LT LV MK RO SI**

(30) Priority: **23.04.1998 US 64623**


(71) Applicant:  
**Newbridge Networks Corporation**  
Kanata, Ontario K2K 2E6 (CA)

(72) Inventors:  
• **Duckering, Brent G.**  
Vancouver, British Columbia V5V 4H4 (CA)  
• **Robotham, Rob E.**  
Ottawa, Ontario K1H 5L8 (CA)

(74) Representative:  
**Hutchins, Michael Richard et al**  
FRY HEATH & SPENCE  
The Old College  
53 High Street  
Horley Surrey RH6 7BN (GB)

(54) **Fair share scheduling of multiple service classes with prioritized shaping**

(57) An apparatus for scheduling multi-service category ATM cell traffic through contention points in the ATM network is provided. The service categories have predefined delivery priorities according to quality of service guarantees. To satisfy these priorities while maintaining fair treatment to low priority connections, aging markers are incrementally assigned to queued cells and these markers in combination with priority data are used to determine which connection is serviced next. A method of scheduling traffic through the network is also disclosed.



**FIGURE 2**

**EP 0 977 402 A3**



European Patent  
Office

# EUROPEAN SEARCH REPORT

Application Number  
EP 99 30 3877

| DOCUMENTS CONSIDERED TO BE RELEVANT  |   |   |  |
|--|---|---|--|
| Category   | Citation of document with indication, where appropriate, of relevant passages   | Relevant to claim   | CLASSIFICATION OF THE APPLICATION (Int. Cl. 5)         |
| X  | PEHA J M: "THE PRIORITY TOKEN BANK: INTEGRATED SCHEDULING AND ADMISSION CONTROL FOR AN INTEGRATED-SERVICES NETWORK"<br>PROCEEDINGS OF THE INTERNATIONAL CONFERENCE ON COMMUNICATIONS (ICC), US, NEW YORK, IEEE,<br>vol. -, 23 May 1993 (1993-05-23), pages 345-351, XP000371116<br>ISBN: 0-7803-0950-2<br>* page 346, right-hand column, line 66 -<br>page 347, left-hand column, line 51 * | 1-4, 13   | H04L12/56<br>H04Q11/04                                 |
| A  | REXFORD J ET AL: "SCALABLE ARCHITECTURES FOR INTEGRATED TRAFFIC SHAPING AND LINK SCHEDULING IN HIGH-SPEED ATM SWITCHES"<br>IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, US, IEEE INC. NEW YORK, vol. 15, no. 5, 1 June 1997 (1997-06-01), pages 938-950, XP000657043<br>ISSN: 0733-8716<br>* page 938, left-hand column, line 8 -<br>line 11 *   | 1-13  | TECHNICAL FIELDS SEARCHED (Int. Cl. 6)<br>H04L<br>H04Q |
| A  | US 5 724 513 A (DE-LEON MOSHE ET AL)<br>3 March 1998 (1998-03-03)<br>* column 4, line 12 *  | 2   |  |
| The present search report has been drawn up for all claims   |   |   |  |
| Place of search<br><b>THE HAGUE</b>  |   | Date of completion of the search<br><b>21 November 2000</b> | Examiner<br><b>Veen, G</b>                             |
| <p><b>CATEGORY OF CITED DOCUMENTS</b></p> <p>X : particularly relevant if taken alone<br/>Y : particularly relevant if combined with another document of the same category<br/>A : technological background<br/>O : non-written disclosure<br/>P : intermediate document</p> <p>T : theory or principle underlying the invention<br/>E : earlier patent document, but published on, or after the filing date<br/>D : document cited in the application<br/>L : document cited for other reasons<br/>.....<br/>&amp; : member of the same patent family, corresponding document</p> |   |   |  |

EPO FORM 1503 03/02 (P4C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT  
ON EUROPEAN PATENT APPLICATION NO.**

EP 99 30 3877

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.  
The members are as contained in the European Patent Office EDP file on  
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

21-11-2000

| Patent document<br>cited in search report | Publication<br>date | Patent family<br>member(s) | Publication<br>date |
|---|---------------------|----------------------------|---------------------|
| US 5724513    A                           | 03-03-1998          | NONE                       |                     |

EPO FORM P4559

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

**THIS PAGE BLANK (USPTO)**